

PATENT APPLICATION

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of

Docket No: CQ10210

Matthew COOPER, et al.

Appln. No.: 10/729,915

Group Art Unit: 2167

Confirmation No.: 4729

Examiner: Robert M. TIMBLIN

Filed: December 9, 2003

For: SYSTEMS AND METHODS FOR ORGANIZING DATA

**DECLARATION OF PRIOR INVENTION TO OVERCOME CITED PATENT OR
PUBLICATION (37 C.F.R. § 1.131)**

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Commissioner:

Purpose of Declaration

1. This declaration is to establish conception of the invention in this application at a date at least prior to August 21, 2003, the effective date of the cited patent application publication:

US2005/0044487 to Bellegarda et al. (application no. 10/644,815).

2. The person making this declaration is one of the joint inventors and is the inventor of the subject matter of the rejected claims. (MPEP 715.04).

Facts and Documentary Evidence

3. To establish the date of conception and diligence from a date prior to the effective date of the reference until the date of filing of the invention of this application, the following attached documents are submitted as evidence:

- FXPAL Internal Invention Proposal signed and witnessed on a date prior to August 21, 2003)

DECLARATION OF PRIOR INVENTION TO OVERCOME
CITED PATENT OR PUBLICATION (37 C.F.R. § 1.131)
U.S. Application No.: 10/729,915

- Email Correspondence regarding preparation of the application between the drafting attorney and the primary inventor (August 7, 2003)
- Email Correspondence regarding a convenient date for inventor interview between the primary inventor and drafting attorney (August 12, 2003)
- Further Email Correspondence regarding a convenient date for inventor interview between the primary inventor and drafting attorney (August 21, 2003)
- Email Correspondence regarding a first draft of the application between the primary inventor and drafting attorney (October 14, 2003)
- Email Correspondence conveying the first draft of the application between the primary inventor and drafting attorney (October 30, 2003)
- Email Correspondence conveying inventor's comments regarding the first draft of the application between the primary inventor and drafting attorney (October 31, 2003)
- Email Correspondence conveying references to the drafting attorney between the primary inventor and drafting attorney (November 7, 2003)
- Email Correspondence conveying another draft to the primary inventor between the primary inventor and drafting attorney (November 13, 2003)
- Email Correspondence verifying inventorship information between the primary inventor and drafting attorney (November 24, 2003)
- Email Correspondence conveying a final draft to the primary inventor between the primary inventor and drafting attorney (December 2, 2003)

DECLARATION OF PRIOR INVENTION TO OVERCOME
CITED PATENT OR PUBLICATION (37 C.F.R. § 1.131)
U.S. Application No.: 10/729,915

4. The attached FXPAL Internal Invention Proposal, has been signed and witnessed on a date prior to August 21, 2003. Thus, the invention in this application has been conceived at least prior to August 21, 2003 which is the effective date of the reference.

Diligence

5. The attached documents also establish diligence beginning from a date prior to the effective date of the reference on August 21, 2003, until the date of filing of the application on December 9, 2003. (MPEP 715.07(a)).

Time of Presentation of the Declaration

6. This declaration is submitted after a final office action together with a Request for Continued Examination (RCE). Therefore, the declaration is timely. (MPEP 715.09).

Declaration

As a person signing below, I hereby declare as follows:

I am a co-inventor of the subject matter claimed in the present patent application. At the time of inventing the subject matter of the present patent application, I was an employee of Fuji Xerox Co. LTD.

On or about a date at least prior to August 21, 2003, I provided the enclosed invention disclosure, directed to the subject matter of the present patent application, to a witness whose signature appears on the disclosure stating that he has read and understood the disclosure.

To my personal knowledge, the copy of the invention disclosure attached to this Declaration is true and correct copy of the invention disclosure that I submitted on or about a date at least prior to August 21, 2003. The copy of the invention disclosure attached to this Declaration has been kept in business records of the law firm in ordinary course of business.

The enclosed invention disclosure establishes that I was in possession of the invention at least prior to August 21, 2003.

DECLARATION OF PRIOR INVENTION TO OVERCOME
CITED PATENT OR PUBLICATION (37 C.F.R. § 1.131)
U.S. Application No.: 10/729,915

Between the date of submission of the invention disclosure and the date of filing of the application on December 9, 2003, I or one of my co-inventors was in contact with the patent attorney drafting the application either directly or through FXPAL patent counsel and diligently reviewed the drafts of the application provided by the drafting attorney.

The enclosed evidence of the communication (10 emails) between the drafting patent attorney and myself, or my co-inventor, establish that I exercised diligence from at least a date prior to the effective date of the reference to the date of filing of the instant application.

The enclosed invention disclosure was created by me and my co-inventors and the draft application submitted by the email of October 30, 2003, was received and reviewed by me or my co-inventor. Further, from my personal knowledge the enclosed invention disclosure and the enclosed draft application are authentic documents that were kept in the business records of the company and/or the law firm in the ordinary course of business.

I hereby declare that all statements made of my own knowledge are true, and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and may jeopardize the validity of the application or any patent issuing thereon.

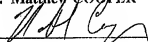
Atty. Docket No.: CQ10210

DECLARATION OF PRIOR INVENTION TO OVERCOME
CITED PATENT OR PUBLICATION (37 C.F.R. § 1.131)
U.S. Application No.: 10/729,915

Signature

A. Inventor(s)

Full name of first inventor: Matthew COOPER

Inventor's signature Date 7/29/08 Country of Citizenship USAResidence Address: 1014 NOB ST SF, CA 94114

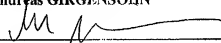
Full name of inventor: Jonathan FOOTE

Inventor's signature _____

Date _____ Country of Citizenship _____

Residence Address: _____

Full name of inventor: Andreas GIRGENSOHN

Inventor's signature Date 7/29/08 Country of Citizenship GermanyResidence Address: 430 Stanford AvePalo Alto, CA 94306

DECLARATION OF PRIOR INVENTION TO OVERCOME
CITED PATENT OR PUBLICATION (37 C.F.R. § 1.131)
U.S. Application No.: 10/729,915

Atty. Docket No.: CQ10210

Signature

A. Inventor(s)

Full name of first inventor: Matthew COOPER

Inventor's signature _____

Date _____ Country of Citizenship _____

Residence Address: _____

Full name of inventor: Jonathan FOOTE

Inventor's signature Jonathan Foote

Date 7/29/2008 Country of Citizenship USA

Residence Address: 755 FLORIDA ST #6

SE CA 94110

Full name of inventor: Andreas GIRGENSOHN

Inventor's signature _____

Date _____ Country of Citizenship _____

Residence Address: _____

FXPAL Invention Proposal FXPAL [REDACTED]

To: Xerox Patent Department
3333 Coyote Hill Rd.
Palo Alto, CA 94304

From: FX Palo Alto Laboratory, Inc.
3400 Hillview Ave., Bldg. 4
Palo Alto, CA 94304

Proposal Submitted By.

Matthew Cooper
PAL 0057
cooper@fxpal.com
(650) 813-7145

Jonathan Foote
PAL0046
foote@fxpal.com
(650) 813-7071

Andreas Girsensohn
PAL 0038
andreasg@fxpal.com
(650) 813-7244

Manager

Lynn Wilcox
PAL 0020
wilcox@fxpal.com
(650) 813-7574

Proposal Title

A Method for Hierarchical Data Clustering

Brief Description

We present a method for hierarchical data clustering which can be applied to a variety of data types. The approach is unsupervised and makes minimal assumptions regarding the structure or statistics of the data. We have applied the technique to event clustering of digital photographs by their timestamps, and present results and comparisons to competing approaches.

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

1 Introduction

This invention is in the field of media analysis. We present a method for hierarchical structural segmentation of media streams, applicable to audio, video, and text. The system has been developed for hierarchical clustering of digital photographs into time-based events. As digital cameras proliferate, the number of digital photographs accumulating on personal computers is growing rapidly. As a result, there is a demand for automatic tools to aid in the organization of these collections.

Digital cameras commonly record a wealth of meta-data into the resulting digital image files. This data is typically stored in a standard Exchangeable Image file format (Exif) [1] image header as part of each JPEG image file. This data includes the timestamp, namely the time the digital photograph was taken, though this information may not be preserved if the image file is subsequently modified. In a recent report, Stanford researchers have found that organizing photos by time significantly improves users' performance in a series of retrieval tasks [2]. Furthermore, consumers often wish to organize their photos in terms of "events" both for browsing and retrieval in their own collections, as well as for sharing selected subsets of photos with others. Events are difficult to define quantitatively or consistently, but most commonly, photographs ascribed to the same event were taken in relatively close proximity in time. Events tend to be difficult to describe in terms of low-level image features, and it is easy to imagine visually dissimilar photos belonging to the same event. For example, pictures from a trip to the beach could include photos, at different times of day, of the sky, beach or ocean, but also of various people.

In such contexts, content-based image similarity may be less useful for photo clustering or event detection than meta-data. In particular, we focus here on processing the photos' timestamps. In the future, many consumer digital cameras will also record GPS location information which should also prove valuable for this task. As a result, we have designed our framework to be flexible to expand in the future to include other meta-data and content-based image features which may prove useful. We also note that there is a wealth of digital file types with timestamps, mainly time of modification as stored by operating system software. These can also be processed similarly to the digital image files. As well, data from system or server logs can be mined for structure using time information using these methods.

2 Technical Details

We illustrate the approach by clustering 512 photos supplied by Fujifilm Software, California (FSOA). All the photos had timestamps and were placed into meaningful folders by the photographer. We process only the timestamps of the photos, and compare our results to the ground truth clustering provided by the folders in which each image was placed. We

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

assume that images in the same folder are part of the same event.

2.1 Pre-processing

Exif headers are processed to extract the timestamp. It should be noted that if the image is modified in several common image viewers and subsequently saved, the Exif information may be lost. In this case, we must rely on the modification time of the digital image file, which will generally be far less reliable for event detection. The first step is to order the photos in time. Denote the number of photos processed as N . If the resulting timestamps are $\{t_i : i = 1, \dots, N\}$, then

$$t_1 \leq t_2 \leq \dots \leq t_N$$

For now, we use minutes as the basic time unit. Note that throughout, we index matrices and novelty-scores by item index (i.e. by time order) and not by absolute time. This is a key difference from previous inventions based on similarity analysis, the time difference between indices (photos) is non-uniform. In fact, time is the (scalar) feature measured for each image, as opposed to more common content-based (vector) features such as color histograms.

2.2 Distance Matrix Embedding

To assess structure in the collection of timestamps, we build a similarity matrix. As an example, the left panel of Figure 1 shows the similarity matrix generated from the ground truth clustering. The elements of the matrix are one for photos from the same folder and zero otherwise. The photos are indexed in time order, such that the (i, j) element of the matrix compares the names of the folders in which i^{th} and j^{th} photos were stored. The blocks along the main diagonal of the matrix are the groups of photos in each folder. A checkerboard pattern along the main diagonal indicates the boundary between folders or events. The crux of the checkerboard pattern is the boundary in time order between the photos in the two events. This is a convenient visualization of the data, which immediately shows that the ground truth event clustering partitions the photos contiguously in time (i.e. events are contiguous and disjoint). For event clustering, we build a family of similarity matrices indexed by the number of photographs according to the general formula

$$S_K(i, j) = \exp\left(-\frac{|t_i - t_j|}{K}\right) \quad (1)$$

The parameter K controls the sensitivity of the exponential similarity measure. K has the unit minutes (the similarity measure is unitless). By varying K , we can visualize the similarity between the timestamps at differing granularities. The left column of Figure 2 shows several similarity matrices computed using (1) for different values of K . As expected, the

Witnessed and Understood by

DBL. CL

Date

1. Submitter Signature

Michael C...

Date

2. Submitter Signature

JP...

Date

3. Submitter Signature

John...

Date

matrices for larger values of K illustrate coarser clusterings of the photos' timestamps. For smaller K the finer dissimilarities between groups of timestamps become apparent. The similarity measure can be tailored to integrate or emphasize other features' such as low-level image features, GPS data, or other meta-data.

2.3 Segmentation

To cluster the photos hierarchically, we compute photo-indexed novelty scores following [3, 4] for each S_K . We take a matched filter approach, correlating a kernel along the main diagonal of S_K . To find boundaries between groups of photos with similar timestamps, we use a Gaussian-tapered 11×11 checkerboard kernel, denoted g and calculate the score

$$\nu_K(i) = \sum_{j,k=-5}^5 S_K(i+j, i+k) g(j, k) \quad (2)$$

For implementation, there is no need to compute the full similarity matrix. Rather, we need only compute the strip around the main diagonal with the same width as the kernel. This reduces computational complexity to become approximately linear with the number of photos, N .

The right panel of Figure 1 shows the novelty score computed from the ground truth similarity matrix using g . Figure 2 shows several similarity matrices and the corresponding novelty scores computed for $K = 10^3, 10^4, 10^5$ minutes. The matrices reveal structure at different resolutions. In turn, the peaks in the corresponding novelty scores indicate a hierarchical set of cluster boundaries between contiguous groups of photos with similar timestamps. In the plots of the right column, the ground truth novelty score of Figure 1 is superimposed in blue.




Figure 1: The left panel shows the ground truth similarity matrix. The right panel shows the novelty score computed using a gaussian checkerboard kernel.

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date




Figure 2: The left panels show the similarity matrices S_K for $K = 10^8$ (top row), $K = 10^6$ (middle row), and $K = 10^5$ (bottom row). The right panels show the corresponding novelty scores computed using a gaussian checkerboard kernel in green, with the ground truth novelty score in blue.

2.4 Clustering via Boundary Detection

To assess performance we locate peaks in the novelty score at each scale, enforcing a hierarchical structure on the detected boundaries. We perform the analysis from coarse scale to fine (decreasing K). To build a hierarchical set of event boundaries, we include the event boundaries detected at coarse scales in the boundary lists for all finer scales. Additional study of the novelty score helps to motivate this approach. Figure 3 shows the

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

novelty scores evaluated at the each of the ground truth cluster boundaries, as they vary with K . Clearly, the novelty scores at these boundary points vary considerably depending on the scale of the analysis. This is produced by the fact that the photos in the clusters on either side of the boundary exhibit different within-class similarity for different values of K per (1). This in turn varies the strength of the correlation with the checkerboard kernel.

The maximum possible novelty score is determined by the maximum of the similarity measure (one, here) and the kernel correlated along the main diagonal of the similarity matrix. For the 11×11 tapered checkerboard kernel employed here, the maximum novelty is 40. Intuitively, Figure 3 shows that different boundaries peak in novelty at different scales. In our application, this stems from the fact that different events have different time extents. To better illustrate this point, Figure 4 shows the boundary at photo index 285 as it appears in two different similarity matrices. Thus different boundaries will be detected for particular values of K in (1).




Figure 3: The figure shows the novelty scores computed for the ground truth cluster boundaries varying with K in (1).

Our approach is based on the assumption that detected event boundaries must, at some scale (i.e. for some K), approach maximum novelty. For each scale K , we detect peaks in the novelty score by analysis of its first difference. Each peak must exceed a threshold of 20. We avoid spurious peaks by using a minimum threshold¹. Specifically, we select one peak in each contiguous region of the novelty score which is greater than 5. The minima in the score correspond to regions of high self-similarity in S (i.e. low novelty). Thus we preferentially locate boundaries between regions of high self-similarity. We analyze the matrices ordered by decreasing value of K , and impose a hierarchical structure on the

¹Both the minimum and maximum thresholds should be determined according to the kernel employed, as described earlier.

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

detected event boundaries.




Figure 4: The left panel shows an excerpt from the similarity matrix computed for $K = 10$. The right panel shows an excerpt from the similarity matrix computed for $K = 1000$. The variability in the similarity within the clusters around the boundary at 285 produce the variations in the novelty score shown in Figure 3.

2.5 Selecting a level of the hierarchy

This approach hierarchically characterizes structure in the photo collection at multiple resolutions. To present the results of the event clustering to users, we automatically rank each level in the hierarchy of detected boundaries. For this we employ a confidence measure based on the average within-class similarity and the between-class dissimilarity of the imagery. Denote the detected boundaries at each level, $B_K = \{b_1, \dots, b_{n_K}\}$. Again, the boundaries, like the similarity matrices, are indexed by photo not by time. For convenience, assume that b_1 and b_{n_K} are the indices of the first (1) and last photos (N), respectively. We then compute the confidence score

$$C(B_K) = \sum_{i=1}^{|B_K|-1} \frac{1}{(b_{i+1} - b_i)^2} \sum_{\substack{j=b_i \\ j \neq i}}^{b_{i+1}} S_K(i, j) - \sum_{i=1}^{|B_K|-2} \frac{1}{(b_{i+1} - b_i)(b_{i+2} - b_{i+1})} \sum_{j=b_i}^{b_{i+1}} \sum_{k=b_{i+1}}^{b_{i+2}} S_K(j, k). \quad (8)$$

The terms in the first sum above quantify the average within-class similarity between the images within each cluster. The second sum quantifies the average between-class similarity between images in adjacent clusters. By negating this sum in the score, we quantify the between-cluster dissimilarity. The confidence measure combines the sum of each cluster's

Witnessed and Understood by

Walter C.

Date

1. Submitter Signature

Walter C.

Date

2. Submitter Signature

John R.

Date

3. Submitter Signature

John R.

Date

average similarity and the sum of the dissimilarity between adjacent clusters. Figure 5 illustrates the idea graphically. The within-class similarity terms are the averages over the terms of darker regions along the main diagonal. The between class terms are the average of the gray regions off the main diagonal.




Figure 5: The figure illustrates the computation of the confidence score. The dark regions represent within-cluster similarity, while the gray regions represent between-cluster similarity.

This confidence measure depends on both the number of detected clusters and K explicitly. Figure 6 illustrates its behavior. The three panels show the regions of the respective similarity matrices averaged and summed to form the confidence measure of (3) for $K = 1778.28$ (top left), $K = 1000$ (top right), and $K = 562.84$ (bottom). For visualization purposes, we have set the remainder of the matrices to zero (elements not contributing to $O(B_K)$). For $K = 1778.28$, we expect a lower confidence score than for $K = 1000$, since the upper right panel shows fewer clusters in number and clustered regions of relatively low similarity. On the other hand, the matrix for $K = 562.84$ shows more clusters than for $K = 1000$, but because K is smaller, again regions of low similarity are clustered. In this manner, the confidence measure emphasizes the most appropriate scale for similarity analysis. The algorithm is summarized below.

Witnessed and Understood by

[Signature]

Date

[Redacted]

1. Submitter Signature

[Signature]

Date

[Redacted]

2. Submitter Signature

[Signature]

Date

[Redacted]

3. Submitter Signature

[Signature]

Date

[Redacted]

Algorithm 1 [Hierarchical Event Detection]

1. Sort photo's timestamps from first to last: $\{t_1, \dots, t_n\}$, $t_i \leq t_j \Leftrightarrow i \leq j$.
2. For each K in decreasing order.
 - (a) Compute the similarity matrix S_K using (1). (Only part of the matrix need be computed).
 - (b) Compute the novelty score v_K of (2).
 - (c) Detect peaks in the novelty score.
 - (d) Form event boundary list using event boundaries from previous iterations and newly detected peaks.
3. Compute confidence score using list of event boundaries, E_K computed for each K following (3).
4. Select event boundary list for K maximizing the confidence score.

2.6 Experimental Results

The precision and recall results for experiments with 1036 images from the FSCA data set appear in Table 1 in the Appendix. The exponential family of similarity measures sweep out a broad range in the precision and recall of the detected cluster boundaries. The confidence score is shown versus the resolution parameter K in Figure 7.

For additional testing, we compare our approach to two other approaches described in the literature [2, 5] and to a simple (constant) threshold. The first two algorithms are closely related. They first sort the photographs in time order and compare the time difference between adjacent photographs to an adaptively computed threshold based on the logarithm of the average inter-photo time difference over a local window of photos. When the time difference exceeds the threshold, an event boundary is labelled between the two photos. For testing, we applied the four algorithms to a collection of 1036 photos taken over 15 months². The photos had time stamps and were placed into meaningful folders by the photographer. We used those folders as ground truth for event boundaries (57 event boundaries). For each algorithm, we measured the precision and recall for the detected event boundaries:

$$\text{Precision} = \frac{\text{Number of correctly detected boundaries}}{\text{Total number of detected boundaries}} \quad (4)$$

$$\text{Recall} = \frac{\text{Number of correctly detected boundaries}}{\text{Total number of true boundaries}} \quad (5)$$

²The 512 photos used to generate the Figures earlier in the paper were the first 512 photos taken in the set of 1036 photos used for testing and comparison against other methods.

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date




Figure 6: The three panels show the regions of the respective similarity matrices averaged and summed to form the confidence measure of (3). The top row shows the matrix for $K = 1778.28$ (left) and $K = 1000$ (right). The bottom panel shows the matrix for $K = 562.34$.

and also calculated the F-score as

$$F\text{-score} = \frac{2 \cdot (\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}} \quad (6)$$

The algorithm described in [5] found all event boundaries but also detected many additional ones (precision 0.39, recall 1.0; F-score 0.584, 0 missed, 88 false positives). Because we wanted to see whether the algorithm described in [2] could produce better results than simple thresholding, we skipped its thresholding step and looked at the first level of the hierarchy it created. At that level, it produced similar results as the algorithm described in [5] (precision 0.38, recall 1.0; F-score 0.548, 0 missed, 94 false positives).

The self-similarity and the simple threshold algorithms have a maximum F-score when both the threshold and K equal 688.33 minutes (11.5 hours). At that value, the self-similarity algorithm has an F-score of 0.833 (precision 0.79, recall 0.88, 7 missed, 13 false positives). The simple threshold algorithm has an F-score of 0.818 (precision 0.72, recall 0.95, 3 missed, 21 false positives). At most time intervals, the self-similarity algorithm performs slightly better than the simple threshold algorithm. Also, the self-similarity

Witnessed and Understood by: John C. [Signature]

Date: [Redacted]

1. Submitter Signature: [Signature]

Date: [Redacted]

2. Submitter Signature: [Signature]

Date: [Redacted]

3. Submitter Signature: [Signature]

Date: [Redacted]




Figure 7: The plot shows the confidence score versus $\log_{10}(X)$ for the 1038 photo FSCA data. The corresponding precision and recall results appear in Table 1.

algorithm does not require that the threshold be determined *a priori*. Furthermore, it supplies a hierarchy of event boundaries. This is desirable as different users will have different preferences for the time extent of the events they define within their collections. Our current interface allows users to explore other levels of organization in the hierarchy.

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

[REDACTED]

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

References

- [1] Digital Still Camera Image File Format Standard. Japan Electronic Industry Development Association, 1998. <http://www.plma.net/standards/t10/PIMA15740/exif.htm>.
- [2] A. Grahani, H. Garcia-Molina, A. Paepcke, and T. Winograd. Time as the essence for Photo Browsing Through Personal Digital Libraries. *Proc. Joint Conf. on Digital Libraries*, 2002.
- [3] J. Foote. Methods for the Automatic Analysis of Music and Audio. FXPAL IP-99-019, 1999.
- [4] J. Foote. Automatic Audio Segmentation using a Measure of Audio Novelty. *Proc. IEEE Intl. Conf. on Multimedia and Expo*, 2000.
- [5] J. Platt, M. Czerwinski, and B. Field. PhotoTOC: Automatic Clustering for Browsing Personal Photographs. Microsoft Research Technical Report MSR-TR-2002-17, 2002.
- [6] M. Slaney, D. Ponceleon, and J. Kaufman. Multimedia Edges: Finding Hierarchy in all Dimensions. *Proc. ACM Conf. on Multimedia*, 2001.
- [7] R. Duda and P. Hart. *Pattern Classification & Scene Analysis*. Wiley-Interscience, 1973.
- [8] A. Loui and A. Savakis. Automatic Image Event Segmentation and Quality Screening for Albuming Applications. *Proc. IEEE Intl. Conf. on Multimedia and Expo*, 2000.

Witnessed and Understood by

1. Submitter Signature

2. Submitter Signature

3. Submitter Signature

Date

Date

Date

Date

6 Appendix

Table 1: Table summarizing precision and recall for detected peaks at multiple scales. Test data is 1098 Images from the FSQA data. The row in bold indicates the resolution which maximizes the confidence score.

K	Precision	Recall	F-score	Confidence Score
0.1	0.76	0.84	0.80	4.88
0.18	0.76	0.84	0.80	5.11
0.32	0.76	0.84	0.80	5.88
0.56	0.76	0.84	0.80	7.08
1	0.76	0.84	0.80	8.4
1.78	0.76	0.84	0.80	10.05
3.16	0.76	0.84	0.80	12.0
5.62	0.76	0.84	0.80	14.24
10	0.76	0.84	0.80	16.67
17.78	0.76	0.84	0.80	19.82
31.62	0.76	0.84	0.80	23.4
56.23	0.76	0.84	0.80	27.24
100	0.76	0.84	0.84	30.03
177.83	0.83	0.84	0.88	31.72
316.23	0.89	0.84	0.89	32.01
562.34	0.89	0.84	0.89	34.01
1000	0.88	0.81	0.81	35.42
1778.28	0.9	0.61	0.72	25.78
3162.3	0.88	0.88	0.88	22.54
5623.4	0.98	0.44	0.60	17.66
10000.00	0.96	0.85	0.91	14.44
17782.79	1.0	0.28	0.41	8.98
31622.78	1.0	0.18	0.30	0
56234.18	1.0	0.06	0.08	0
100000.00	1.0	0.04	0.07	0
177827.94	1.0	0	0	0
316227.77	1.0	0	0	0
562341.39	-	0	0	0
1000000.00	-	0	0	0

Witnessed and Understood by

Date

1. Submitter Signature

Date

2. Submitter Signature


Date

3. Submitter Signature

Date

Communications Center

JAO/ssk/File

From: Roe, Stephen J. 
Sent: Thursday, August 07, 2003 4:39 PM
To: cooper@fxpal.com
Cc: Kim, Seth S.
Subject: (Archive Copy) FX/A2018 (O&B Ref.)

PRIVILEGED AND/OR CONFIDENTIAL:

This e-mail (including any attachments) is intended only for the use of the individual or entity named above and may contain privileged or confidential information. If you are not the intended recipient, or the employee or agent responsible to deliver it to the recipient, you are notified that any review, dissemination, distribution or copying of this e-mail is prohibited. If you have received this e-mail in error, please immediately notify us by e-mail, facsimile, or telephone; return the e-mail to us at the e-mail address below; and destroy all paper and electronic copies.

Communications Center
Oliff & Berridge, PLC
277 S. Washington Street, Suite 500
Alexandria, VA 22314
Tel: 703-836-6400
Fax: 703-836-2787
Email: commcenter@oliff.com

MESSAGE:

Re: New U.S. Patent Application
Inventor: Matthew COOPER et al.
Our Ref: 114891
Your Ref: FX/A2018

Dear Matt:

We have been assigned by PARC to prepare another new US patent application based on one of your invention proposals for FXPAL, in particular, for your invention titled "A method of hierarchical data clustering." We would like to arrange a telephone conference call at a time mutually agreeable to you and your co-inventors, Jonathan Foote and Andreas Girgensohn, to discuss the invention proposal, including the current status of the development of the invention. We expect the telephone conference to last up to an hour to discuss all aspects of the invention.

Please call or email either of us to discuss the time for a telephone conference.

Regards,

Steve Roe and Seth Kim

Communications Center

SSK

From: M. Cooper [Cooper@fxpal.com]
Sent: Tuesday, August 12, 2003 1:25 PM
To: Roe, Stephen
Cc: Kim, Seth S.; Jon Foote
Subject: (Archive Copy) Re: FX/A2018 (O&B Ref. 114891)

Hi Steve and Seth,

[REDACTED] I'm hoping we can schedule the
teleconference for sometime next week. Thursday morning (pst) is a good
time for me.

Thanks,
Matt Cooper

Roe, Stephen J. wrote:

>
> PRIVILEGED AND/OR CONFIDENTIAL:

> This e-mail (including any attachments) is intended only for the use
> of the individual or entity named above and may contain privileged or
> confidential information. If you are not the intended recipient, or
> the employee or agent responsible to deliver it to the intended
> recipient, you are notified that any review, dissemination,
> distribution or copying of this e-mail is prohibited. If you have
> received this e-mail in error, please immediately notify us by e-mail,
> facsimile, or telephone; return the e-mail to us at the e-mail address
> below; and destroy all paper and electronic copies.

> Communications Center
> Oliff & Berridge, PLC
> 277 S. Washington Street, Suite 500
> Alexandria, VA 22314
> Tel: 703-836-6400
> Fax: 703-836-2787
> Email: commcenter@oliff.com <mailto:commcenter@oliff.com>

> -----
> MESSAGE:

> -----
> Re: New U.S. Patent Application
> Inventor: Matthew COOPER et al.
> Our Ref.: 114891
> Your Ref.: FX/A2018

> Dear Matt:

>
> We have been assigned by PARC to prepare another new US patent
> application based on one of your invention proposals for FXPAL, in
> particular, for your invention titled "A method of hierarchical data
> clustering." We would like to arrange a telephone conference call at
> a time mutually agreeable to you and your co-inventors, Jonathan Foote
> and Andreas Girgensohn, to discuss the invention proposal, including
> the current status of the development of the invention. We expect the
> telephone conference to last up to an hour to discuss all aspects of
> the invention.

>
> Please call or email either of us to discuss the time for a telephone
> conference.

Kim, Seth S.

From: Roe, Stephen J.
Sent: Wednesday, August 20, 2003 5:35 PM
To: 'Matt Cooper'
Cc: Kim, Seth S.
Subject: RE: FX/A2018 (O&B Ref. 114891)

Matt:

[REDACTED]
tomorrow or Friday are fine. We assume by tomorrow morning you mean 9-12 PDT. If so, we are available from 10am-12 noon tomorrow morning PDT (1-3 pm our time). If that is convenient, please send me a confirming email, and confirm your telephone number.

We look forward to discussing this application with you.

Regards,

Steve Roe and Seth Kim

-----Original Message-----

From: Matt Cooper [mailto:cooper@fxpal.com]
Sent: Wednesday, August 20, 2003 5:21 PM
To: Roe, Stephen J.
Subject: Re: FX/A2018 (O&B Ref. 114891)

Hi Stephen,

[REDACTED]
I'm free this week and next to schedule a phone conference. Tomorrow morning or Friday afternoon is fine, so let me know what's convenient for you.

Thanks,
Matt Cooper

Roe, Stephen J. wrote:

>
> PRIVILEGED AND/OR CONFIDENTIAL:
>
> This e-mail (including any attachments) is intended only for the use
> of the individual or entity named above and may contain privileged or
> confidential information. If you are not the intended recipient, or
> the employee or agent responsible to deliver it to the intended
> recipient, you are notified that any review, dissemination,
> distribution or copying of this e-mail is prohibited. If you have
> received this e-mail in error, please immediately notify us by e-mail,
> facsimile, or telephone; return the e-mail to us at the e-mail address
> below; and destroy all paper and electronic copies.
>

> Communications Center
> Oliff & Berridge, PLC
> 277 S. Washington Street, Suite 500
> Alexandria, VA 22314
> Tel: 703-836-6400
> Fax: 703-836-2787
> Email: commcenter@oliff.com <mailto:commcenter@oliff.com>
> -----

> MESSAGE:

Communications Center

From: Matt Cooper [cooper@fxpal.com]
Sent: Thursday, August 21, 2003 2:50 PM
To: Roe, Stephen J.
Subject: (Archive Copy) Re: FX/A2018 (O&B Ref. 114891)

ICME00_jou_savaki s.pdf #3
jtf_icme.pdf #4
MSR-tr-2002-17.pdf #5
Slaney2001.pdf #6
stanford.pdf #2

I attached several of the references. Reference [3] is issued patent 6,542,869.

thanks,
Matt

Roe, Stephen J. wrote:

>
>
> PRIVILEGED AND/OR CONFIDENTIAL:

>
> This e-mail (including any attachments) is intended only for the use
> of the individual or entity named above and may contain privileged or
> confidential information. If you are not the intended recipient, or
> the employee or agent responsible to deliver it to the intended
> recipient, you are notified that any review, dissemination,
> distribution or copying of this e-mail is prohibited. If you have
> received this e-mail in error, please immediately notify us by e-mail,
> facsimile, or telephone; return the e-mail to us at the e-mail address
> below; and destroy all paper and electronic copies.

>
> Communications Center
> Oliff & Berridge, PLC
> 277 S. Washington Street, Suite 500
> Alexandria, VA 22314
> Tel: 703-836-6400
> Fax: 703-836-2787
> Email: commcenter@oliff.com

> -----
> MESSAGE:

> -----
> Matt:

>
> We will call you at 10am your time today. We look forward to going
> over your IP with you.

>
> Regards,

>
> Steve Roe and Seth Kim

>
> -----Original Message-----

> From: Matt Cooper [mailto:cooper@fxpal.com]
> Sent: Wednesday, August 20, 2003 6:15 PM
> To: Roe, Stephen J.
> Subject: Re: FX/A2018 (O&B Ref. 114891)

>
>
> Stephen,

Kim, Seth S.

From: Roe, Stephen J.
Sent: Tuesday, October 14, 2003 2:16 PM
To: "M. Cooper"
Cc: Kim, Seth S.
Subject: RE: FX/A2018 (O&B Ref. 114891)

Matt:

We are finalizing the first draft and expect to get it to you this Friday. Will that work with your schedule?

[REDACTED]

Thanks,

Steve Roe and Seth Kim

-----Original Message-----

From: M. Cooper [mailto:mcooper@fxpal.com]
Sent: Tuesday, October 14, 2003 2:11 PM
To: Roe, Stephen J.
Subject: Re: FX/A2018 (O&B Ref. 114891)

Hi Steve,

Did we get this application out?

[REDACTED]

Thanks,
Matt

Roe, Stephen J. wrote:

>
>
> PRIVILEGED AND/OR CONFIDENTIAL:
>
> This e-mail (including any attachments) is intended only for the use
> of the individual or entity named above and may contain privileged or
> confidential information. If you are not the intended recipient, or
> the employee or agent responsible to deliver it to the intended
> recipient, you are notified that any review, dissemination,
> distribution or copying of this e-mail is prohibited. If you have
> received this e-mail in error, please immediately notify us by e-mail,
> facsimile, or telephone; return the e-mail to us at the e-mail address
> below; and destroy all paper and electronic copies.

>
> Communications Center
> Oliff & Berridge, PLC
> 277 S. Washington Street, Suite 500
> Alexandria, VA 22314
> Tel: 703-836-6400
> Fax: 703-836-2787
> Email: commcenter@oliff.com

> -----
> MESSAGE:
> -----

Kim, Seth S.

From: Roe, Stephen J.
Sent: Thursday, October 30, 2003 7:26 PM
To: 'cooper@fxpal.com'
Cc: Kim, Seth S.
Subject: FW: Your Ref.: FX PAL IP-02-018; Our Ref.: 114891

Re:	New U.S. Patent Application
	Inventor: Matthew COOPER et al.
	Our Ref.: 114891
	Your Ref.: FX PAL IP-02-018

Dear Matt:

Further to our telephone conference earlier today, we enclose a first draft of your above-identified patent application in electronic form in MS Word. Further claims will be forwarded to you by tomorrow morning. The drawings are also being sent by facsimile to 650-813-7081 concurrently with this e-mail message.

Steve Roe and Seth Kim



Draft application
letter.pdf



First Draft
application FX-A20...

Communications Center

SSK
87

From: Matthew Cooper [cooper@fxpal.com]
Sent: Friday, October 31, 2003 11:35 AM
To: Roe, Stephen J.
Cc: Kim, Seth S.; sumiya@fxpal.com
Subject: (Archive Copy) Re: FW: Your Ref.: FX PAL IP-02-018; Our Ref.: 114891



114891.mlc
revised.doc

Stephen and Seth,

I'm attaching my revisions which were few; the application looks generally good.

please call me with any questions at 650 813 7145.

Thanks,
Matt Cooper

Roe, Stephen J. wrote:

>
> PRIVILEGED AND/OR CONFIDENTIAL:
>
> This e-mail (including any attachments) is intended only for the use
> of the individual or entity named above and may contain privileged or
> confidential information. If you are not the intended recipient, or
> the employee or agent responsible to deliver it to the intended
> recipient, you are notified that any review, dissemination,
> distribution or copying of this e-mail is prohibited. If you have
> received this e-mail in error, please immediately notify us by e-mail,
> facsimile, or telephone; return the e-mail to us at the e-mail address
> below; and destroy all paper and electronic copies.
>
> Communications Center
> Oliff & Berridge, PLC
> 277 S. Washington Street, Suite 500
> Alexandria, VA 22314
> Tel: 703-836-6400
> Fax: 703-836-2787
> Email: commcenter@oliff.com <mailto:commcenter@oliff.com>

Communications Center

From: M. Cooper [cooper@fxpal.com]
Sent: Friday, November 07, 2003 7:30 PM
To: Kim, Seth S.
Cc: foote@fxpal.com; Roe, Stephen J.
Subject: (Archive Copy) references



heckerman96tutorial.pdf
l.pdf

Seth,

Here are references for the BIC (Bayes information criterion). The first one is a general review of the BIC, Bayesian learning, and other

David Heckerman. A tutorial on learning with bayesian networks.
Technical Report MSR-TR-95-06, Microsoft Research, Redmond, Washington,
1995. Revised June 96. <http://citeseer.nj.nec.com/heckerman96tutorial.html>

S. Chen and P. Gopalakrishnan, "Speaker, environment and channel change detection and clustering via the Bayesian information criterion," in DARPA speech recognition workshop, 1998.
<http://citeseer.nj.nec.com/chen98speaker.html>

S. Renals and D. Ellis (2003). Audio Information Access from Meeting Rooms <<http://www.ee.columbia.edu/~t7Edpwe/pubs/icassp03-mtg.pdf>>
/Proc. ICASSP-03/, Hong Kong, April 2003

thanks for your help, let me know if you need anything else.

-- INC

Matthew Cooper
cooper@fxpal.com <mailto:cooper@fxpal.com>

2000-01-01

Kim, Seth S.

From: Roe, Stephen J.
Sent: Thursday, November 13, 2003 11:15 AM
To: 'cooper@fxpal.com'
Cc: Kim, Seth S.
Subject: FX/A2018 (O&B Ref. 114891) draft non-provisional application

Re:	New U.S. Patent Application
	Inventor: Matthew COOPER et al.
	Our Ref.: 114891
	Your Ref.: FX/A2018

'Matt:

We have now completed, and enclose, a draft of the non-provisional application as promised, along with the formal figures and a cover letter. Please review the cover letter before you review the application.

Thanks,

Steve Roe and Seth Kim



Draft Application
Cover Letter...



Draft
1-Provisional Applica



FIGURES.pdf

Kim, Seth S.

From: Kim, Seth S.
Sent: Monday, November 24, 2003 11:48 AM
To: Roe, Stephen J.
Subject: Your Ref.: FX PAL IP-02-018; Our Ref.: 114891

Please forward to: cooper@fxpal.com

Re:	New U.S. Patent Application
	Inventor: Matthew COOPER et al.
	Our Ref.: 114891
	Your Ref.: FX PAL IP-02-018

Dear Matt:

We have the following information for you and your co-inventors, Jonathan Foote and Andreas Girgensohn. Please confirm their accuracy and notify us if any of the information should be updated.

Mathew L. Cooper
3998 23rd Street.
San Francisco, CA 94114 USA
Citizenship: USA

Andreas Girgensohn
210 Waverly Street #6
Menlo Park, California 94025 USA
Citizenship: Germany

Jonathan T. Foote
450 Laurel Street.
Menlo Park, CA 94025 USA
Citizenship: USA

Very truly,

Steve and Seth

Kim, Seth S.

From: Roe, Stephen J.
Sent: Tuesday, December 02, 2003 3:18 PM
To: Kim, Seth S.
Subject: FW: FX/A2018 (O&B Ref.: 114891)

-----Original Message-----

From: Roe, Stephen J.
Sent: Tuesday, November 25, 2003 4:49 PM
To: 'cooper@fxpal.com'
Subject: FX/A2018 (O&B Ref.: 114891)

Re:	New U.S. Patent Application
	Inventor: Matthew COOPER et al.
	Our Ref.: 114891
	Your Ref.: FX/A2018

Dear Matt:

Thank you for your November 24 e-mail messages. We have attached clean and marked-up versions of the final draft of your above-identified patent application and the formal drawings, Figs. 1-18, for you and your co-inventors' review, along with the Declaration/Power of Attorney and the Assignment. [REDACTED]

Please call either of us if you have any questions.

Regards,

Steve Roe and Seth Kim



Final Draft
Application.zip